Contribution to the DOE National Collaboratories Program Investigator's Meeting, August 10-12 at ANL.

Cyber-Security Priorities: Building Secure Collaboratories

Frank Siebenlist (franks@mcs.anl.gov), Von Welch (vwelch@ncsa.uiuc.edu), Sam Meder (meder@mcs.anl.gov)

Abstract: This paper discusses the need for a Grid authorization policy architecture required to support the structure of distributed, multi-organizational collaborations. This architecture aims to address the challenges of authorization for multi-domain scientific collaborations, while it should address issues of accommodating a diverse set of policy sources and mechanisms in order to dynamically produce authorization policy which allows for controlled resource sharing while maintaining the security requirements of all parties involved (collaborations, sites, users, and oversight bodies). Implementations of this architecture should integrate authorization related services, like VOMS, SAZ, PERMIS and CAS, and should allow for the deployment of rich policy expression languages, like XrML or XACML.

Scientific advances today are almost exclusively the result of large collaborative team efforts. The Department of Energy SciDAC program contains many examples of such collaborative teams, like *The Particle Physics Data Grid* (PPDG), *The Earth Systems Grid* (ESG), and *The National Fusion Collaboratory* (NFC). All these project share the following four essential properties of collaborative work:

- *Geographical and Organizational Distribution*. Participants in a collaborative activity are distributed, both geographically and organizationally, as are the tools and resources used to perform the work of the collaboration (e.g., computers, data sets, storage devices, simulation programs).
- *Large and Dynamic Scale*. Collaborations can scale in size from a few individuals to literally hundreds or thousands of participants, which is the case of many high-energy physics experiments. Furthermore, the membership of participants of a collaboration is not static, but frequently varies over the lifetime of the collaborative task. Participants may join or leave, and resources may be added or removed
- *Diverse Roles*. Collaborations may span areas of expertise, with members filling different roles within the collaboration. The role of a member may be fixed for the duration of the collaboration, or it may change during its lifetime. For example, in the case of climate modeling there are distinct roles with associated specific rights of the simulation writer, the simulation runner, and the consumer of simulation data.
- *Community resources*. The work of the team is enabled by providing team members with access to a variety of resources including computers, storage systems, datasets, applications, and tools. Thus in a real sense, a collaboration is not just the group of individuals participating in the activity, but the resources that can be used by members of the collaboration to conduct their work.

In order for such a collaboration to succeed, its participants must have means to perform the work of the collaboration: e.g., mechanisms for annotating and cataloging information so that it may be understood by members of the collaboration, electronic notebooks for sharing what processes had been followed, interfaces that make computing resources available for use, methods for discovering and initiating simulation codes, etc. Most work on collaboration tools focuses on the development of these tools.

In order for the collaboration to be effective, its participants must also have mechanisms for establishing and maintaining the structure of the collaboration. Users and resources are rarely fully devoted to any particular collaboration, but instead remain bound by policies and technologies that are in place at their home organization and in other collaborations of which they are a part.

The foundation of these mechanisms for collaboration structure is the ability to identify the collaboration members (authentication) and determine what their specific roles and privileges are (authorization). This process of authentication and authorization is non-trivial due to the number of parties that contribute policy to authorization process – the collaboration, the resource sites, the resource managers, and the users are all examples of entities, which have a stake in the authorization process.

The requirement that the applied policy in a collaboratory be based on a combination of policies from a number of sources:

- The resource owner: the owner will need to specify how the collaboration may go about using the resource. Since resources are often selectively shared with the collaboration and not given over completely, the owner needs to be able to maintain ultimate control of the resource and needs the ability to express the subset of rights or control it is willing to delegate to the collaboration. This includes potentially specifying how much of the resource they may use (e.g., amount of storage, bandwidth or compute cycles), when they may use the resource, what sorts of members the collaboration may allow access (e.g., members may have had to agree to an policy of acceptable use), and to what degree they may use the resource (e.g., read versus write, able to install their own applications or just used those provided by the owner).
- The site: the site hosting resources may have policies that typically must be applied ubiquitously across all resources at a site, such as auditing, acceptable use polices, etc.
- The collaboration: the collaboration needs to specify its membership, the roles of those members and what privileges those roles convey (e.g., how much of each type of resource a user in a particular role can consume, what data they can access). The collaboration may also need to specify their own policies in regards to the coordination of the contributed resources (e.g., no user may use more than a given fraction of the total resources at a given moment; all resource usage can be preempted by a subset of the collaboration if needed for a critical task).
- The collaboration users: users will need to delegate rights to processes and other users to act on their behalf. For example, a computational job may need to access the user's data or resources for storage; a colleague may need access to an online document for a short-term collaboration.
- Other stake holders: resource owners, the collaboration and the users involved may also be members of larger bodies that dictate policies that effect the collaboration. For example, DOE laboratories contributing resources to collaborations must abide by DOE policies regarding auditing, authentication, etc.

The complexity of combining these different policies is increased by several factors:

The policies to be combined may be dynamic and diverse. The collaboration may need to change its policy as users or contributed resources change, when its goals change, or the collaboration may begin with only a rudimentary idea of what its policies should be and formulate more appropriate policies over time. Policies from the site or other parties may similarly change. This requires mechanisms for forming the combined policy must be capable of acquiring and combining these policies in real-time as the policy is applied.

Policies and the attribute and identity credentials they depend on, come in a variety of forms. Unfortunately we have many standards for expressing policies and attributes. Different sites and collaborations have in the past, and will continue in the future, continue to select different mechanisms (e.g. X.509 attribute certificates, SAML assertions, LDAP) for legitimate reasons. A system that supports the combining of policies and attributes from multiple sources must be capable of dealing with different formats.

Policies vary in granularity. In some cases policy will need to be fine-grained, e.g., expressing access to individual files, while in other cases they will be rather coarse-grained, e.g., only users with DOE Grids identity certificates may use this service. This requires that the enforcement system for the resulting policy be flexible. For example, a system designed to only enforce access control to a site may not be capable of differentiating between what file is being referenced in a request

Policies cover a wide range of resources. Collaborations are not concerned with only a single type of resource (e.g., computation, data storage, visualization, instruments), but with a wide range. This requires the policy and enforcement systems to be broadly applicable.

As such, we've identified the fundamental issue of enabling the combining of these different policies to allow for structured collaborations, and believe that there is an urgent need to develop scalable, secure, and usable methods and tools for combining policies in collaboration from all participants in the collaboration, including sites, users, the collaboration itself and other stakeholders. We also would like to stress that any implementations of this architecture should integrate with ongoing efforts for authorization related services, like VOMS, SAZ, PERMIS and CAS, and should allow for the deployment of rich policy expression languages, like XrML or XACML.